# LOCAL RECONFIGURATION OF ATM VIRTUAL PATH CONNECTION NETWORKS

S.A. Berezner<sup>1</sup>, J.M. de Kock<sup>2</sup>, A.E. Krzesinski<sup>2</sup> and P.G. Taylor<sup>3</sup> \*

<sup>1</sup>Dept of Maths & Stats University of Melbourne Victoria 3052 Australia berezner @ms.unimelb.edu.au <sup>2</sup>Dept of Computer Science University of Stellenbosch Stellenbosch 7600 South Africa { jmdekock, aek1 } @cs.sun.ac.za <sup>3</sup>Dept of Applied Maths University of Adelaide Adelaide 5005 Australia ptaylor @maths.adelaide.edu.au

Abstract Modern communication networks based on SDH/SONET or ATM can construct logical end-to-end connections between all origin-destination (O-D) pairs and thus create fully meshed logical networks upon sparse physical networks. Such logical connections are known as virtual path connections (VPCs) and a logical network formed by VPCs is known as a VPC network (VPCN). We have developed an efficient algorithm called XFG to compute optimal VPCNs for carrying multiservice traffics. The XFG algorithm needs to know the call arrival intensities between all O-D pairs. These data may not always be available. We have therefore developed a local reconfiguration algorithm to augment the XFG algorithm. The local algorithm makes use of data which are likely to be available in real networks. A call arrival at an O-D pair potentially triggers a local reconfiguration which affects only those routes that either connect the O-D pair or which use the physical link which connects the O-D pair. The information necessary to calculate the reconfiguration is kept locally at the O-D pair.

We present simulation experiments which apply local reconfiguration to a model of a small ATM network. The initial experiments indicate that the local reconfiguration algorithm performs well on the network model under investigation.

**Keywords:** alternative routing; ATM networks; dynamic reconfiguration; non-linear optimization; virtual path connection networks.

\*The second and third authors of this work were supported by the the South African National Research Foundation Grant No. 2034100, Telkom SA Limited and Siemens Telecommunications. The fourth author was supported by the Australian Research Council Grant No. A69702505.

#### 1. INTRODUCTION

Virtual paths (VPs) and virtual path connections (VPCs) are important concepts in ATM networks. A VP recognizes the distinct identity of a traffic stream between two communicating nodes. A VPC consists of a series of concatenated VPs and specifies a route to be traversed by a set of virtual channel connections from an originating node through a number of intermediate nodes to a destination node. Bandwidth is logically assigned to a VPC by reserving a certain part of the transmission capacity on each physical link for the exclusive use of the VPC. Using VPCs allows for faster setup of new connections along predefined routes and rapid movement of ATM cells with minimal processing at intermediate nodes.

A logical network formed by VPCs is known as a virtual path connection network (VPCN). The VPCN plays a central role in dynamic network re-configuration (see [1] and the references therein) which has been proposed as a simple and robust resource management control to manage ATM networks to respond optimally to traffic variations where these variations occur on a time scale of the order of hours. The network management control measures the offered traffic and, should a significant and persistent traffic change be detected, the control computes and implements an optimal VPCN which maximizes network profit for the new traffics. The motivation is that a dynamic VPCN will be a better method to handle slowly varying traffics, while short term traffic variations due to random mismatches between offered traffics and available bandwidth are best handled by alternative routing.

However, an excessive use of alternative routing may overload the signalling network. It is desirable to limit the number of alternatively routed calls which might arise when the traffic pattern differs from the traffic pattern which was used to design the optimal VPCN. We propose to do this by local reconfiguration: if the proportion of rejected calls for a specific origin-destination (O-D) pair exceeds the GoS requirements, the capacity of the VPC corresponding to this O-D pair is adjusted.

## 2. **RECONFIGURATION ALGORITHMS**

Anerousis and Lazar [1] survey a number of algorithms for constructing VPCs in broadband networks. They provide a taxonomy for the algorithms, distinguishing the algorithms according to a number of criteria, among them whether they are synchronous or asynchronous, whether they are centralized or decentralized, the form of the cost function being optimized, and whether signalling costs are taken into account. Synchronous algorithms update the capacity allocated to a VPC based on observed demand for call establishment. Asynchronous algorithms maintain a capacity allocation for a fixed time. Centralized algorithms are updated at one location and require the collection of up-to-date information from all network nodes, while decentralized algorithms are executed in every switch using only local information.

Another criterion which can also be added is whether an algorithm is data-driven or data-free. A data-driven algorithm uses information about the statistics of the call arrival and service processes. A data-free algorithm does not use information about the processes which load the network: it uses information only about the network states which arise as a result of the traffic processes.

#### 3. GLOBAL RECONFIGURATION

Consider a network which carries calls that belong to S service classes. Let  $b_s$  denote the effective bandwidth of a call of class s where  $1 \leq s \leq S$ . For each O-D pair (i, j) let  $C_{ij}$  denote the capacity of the link i - j,  $\lambda_{ij}^s$ the Poisson arrival rate of class s calls and  $1/\mu_{ij}^s$  the mean holding time of class s calls offered to the O-D pair (i, j).  $C_{ij}$  is zero if the physical link i - j does not exist. Let  $\rho_{ij}^s = \lambda_{ij}^s/\mu_{ij}^s$  denote the arrival intensity of calls of class s offered to the O-D pair (i, j). Let  $\theta_{ij}^s$  denote the rate at which a class s call offered to the O-D pair (i, j) generates revenue. The revenue rate does not depend on how the call is routed although this assumption can be relaxed if cross-connection costs are taken into account. The effective bandwidth  $b_s$  of a call of class s is also independent of the call's route.

Reconfiguration is performed by cross-connecting circuits using configurable switches in transit nodes to form logical direct links (VPCs) between all O-D pairs. A route is a sequence of physical links connecting an O-D pair. A route between an O-D pair (i, j) is active if circuits are reserved on each link of the route for the sole use of calls between iand j. Let  $\mathcal{R}_{ij}$  denote the set of active routes between the O-D pair (i, j). We assume that the physical links are dimensioned to carry their own offered traffics so that for each O-D pair (i, j) connected by a physical link,  $\mathcal{R}_{ij} = \{i - j\}$ . We restrict ourselves to the design of service integrated VPCNs where each VPC carries all traffic classes.

Let  $\mathcal{R} = \bigcup_{ij} \mathcal{R}_{ij}$  denote the set of active routes. Let  $\mathbf{x} = (x_1 x_2 \dots x_R)$ denote the transmission capacity assignments to all active routes in the network where  $x_r$  is the capacity of route r and  $R = |\mathcal{R}|$ . The (logical) capacity of the VPC connecting the O-D pair (i, j) is given by  $\widehat{C}_{ij}(\mathbf{x}) = \sum_{r \in \mathcal{R}_{ij}} x_r$ . Let  $B_{ij}^s(\mathbf{x})$  denote the blocking probability experienced by a class s calls on the VPC connecting the O-D pair (i, j). The VPCN design problem can be specified in terms of the following constrained non-linear optimization problem: Maximize

$$F(\mathbf{x}) = \sum_{ijs} \theta_{ij}^s \rho_{ij}^s (1 - B_{ij}^s(\mathbf{x}))$$
(1)

Subject to:

$$\sum_{r \in \mathcal{A}_{ij}} x_r \le C_{ij} \qquad x_r \ge 0.$$

where  $\mathcal{A}_{ij}$  denotes the set of active routes that use the physical link i-j. Note that  $\mathcal{A}_{ij}$  includes the single-link route i-j.

An obvious way to obtain a solution for the optimization problem 1 is to use constrained non-linear programming (NLP) methods. However, for networks of realistic size, the number of decision variables in the NLP and the computation time required to solve the NLP grow to an extent that this becomes impossible in practice.

We have therefore developed an efficient deterministic algorithm named XFG [2] to solve the optimization problem 1. The XFG algorithm computes an optimal VPCN in a sequence of steps. At each step the algorithm either adds or removes capacity to or from a route r in  $\mathcal{R}_{ij}$ . The route r is chosen so that the increase or decrease of capacity on route r, which is done by seizing or releasing capacity from or to the links along route r, leads to the largest increase in the rate of earning revenue. The XFG algorithm terminates when no such route r can be found. Since at each step the XFG algorithm obtains the largest possible increase in revenue rate and does not take the long-term impacts of the capacity changes performed at each step into account, the XFG algorithm is a one-step greedy algorithm. According to the classification presented above, XFG is an asynchronous, centralized algorithm which maximizes the network revenue rate and does not take signalling costs into account.

The main advantage of the XFG algorithm is its computational efficiency. For example the non-linear optimizer CFSQP [4] requires several hours to compute an optimal VPCN for the model of the NSF ATM backbone presented in section 6 The XFG algorithm solves the same model in a few seconds. Both execution times are for a Pentium II 400 MHz machine.

### 4. LOCAL RECONFIGURATION

The XFG algorithm constructs a fully-meshed VPCN from a (possibly sparse) physical network. Since is is likely that the traffic profiles are continuously changing, the VPCN will not remain optimal forever. The cost of installing an optimal VPCN design is a computationally expensive process for a large network. Global reconfiguration therefore cannot be performed each time the traffic offered to some O-D pair can no longer be carried to the agreed GoS.

We therefore distinguish between two VPCN design problems. A **global** re-design occurs when either the offered traffics or the physical network change significantly. In this case both the composition and the capacities of the the optimal VPCs are recalculated. Traffic initiated global re-designs should occur relatively infrequently, perhaps every hour or so. A **local** re-design occurs in response to a slow though persistent change in the offered traffic intensities. New routes are not created: instead, the capacities of the existing routes are adjusted to afford better service to the changed traffic. Local re-designs may occur every few minutes. Finally, dynamic alternate routing (DAR) [3] is used to deal with the short time-scale random mismatches between offered traffic and VPC capacity.

The effect of local reconfiguration is to make the VPCN act as an "intelligent entity" which can move logical capacity to where is it most needed: if a logical link notices that it is getting close to being over-loaded, then it augments its capacity by relocating capacity from "nearby" links. Such a strategy has several desirable features: (1) The strategy is local in nature. Each logical link (i, j) needs to know the offered traffic  $\rho_{ij}$  and the current state of the physical links through which it passes. Each physical link needs to know the current state of the logical links which pass through it. An important consequence of this property is that the strategy is **scalable**: it works independently of the size of the network. (2) The strategy is dynamic in that it is able to respond to changes in traffic patterns.

Local reconfiguration is performed at an O-D pair (i, j) in response to some trigger: for example a reconfiguration can be attempted if the blocking probability  $B_{ij}^s$  for the O-D pair (i, j) violates an agreed GoS. Local reconfiguration will not be performed each time a call is blocked since this would give rise to an excessive signalling overhead. Such a call can either be queued, alternatively routed or lost.

# 5. DATA-FREE LOCAL RECONFIGURATION

The local reconfiguration process described above is data-driven in that it needs to know accurate values for the traffic arrival intensities. These may not be known for the following reasons: they are not measured by the network management system; they are measured but, to reduce overhead costs, they are measured infrequently; competing service providers operate separate logical networks on the same physical infra-structure – it is likely that the competitors will be denied access to each others traffic data.

We first present an abuse of notation which facilitates the description of the data-free local reconfiguration algorithm. Consider the routes that use the physical link i - j. One of these routes is the single-link route twhich carries the traffic offered to the O-D pair (i, j). Let  $\mathcal{T} \subset \mathcal{R}$  denote the set of all single-link routes. Consider a multi-link route r. With an abuse of notation let  $t \in r$  denote the single-link route t where  $\mathcal{R}_{ij} = \{t\}$ and  $\{i - j\} \in r, t \in \mathcal{T}$  and  $r \in \mathcal{R}$ .

The data-free local reconfiguration algorithm assumes the existence of a VPCN and that the following information is available:

- 1. For each physical link i j the algorithm knows: the set of routes  $\mathcal{A}_{ij}$  which use the link i j, the allocated capacity  $x_r$  and the unused capacity  $y_r$  on each route  $r \in \mathcal{A}_{ij}$  where  $y_r = x_r \sum_s n_{rs} b_s$  where  $n_{rs}$  is the number of class s calls in service on route r and  $b_s$  is the effective bandwidth of a call of class s.
- 2. For each O-D pair (i, j) not connected by a physical link the algorithm knows: the set of routes  $\mathcal{R}_{ij}$  which connect the O-D pair (i, j), the capacity  $x_r$  allocated to each route  $r \in \mathcal{R}_{ij}$ , and the unused capacity  $y_t$  on each single-link route  $t \in r$ .

Let  $X_{ij} = \sum_{r \in \mathcal{R}_{ij}} (x_r - \sum_s n_{rs} b_s)$  denote the free capacity of the VPC connecting the O-D pair (i, j). If the free capacity  $X_{ij}$  falls below a threshold  $D_{ij}$  then the local reconfiguration algorithm will attempt to add U units of capacity to  $X_{ij}$ . There are two possibilities:

- 1. The physical link i j exists: Select one of the longest routes  $r \in \mathcal{A}_{ij}$  with unused capacity  $y_r > U$ . Decrease the capacity  $x_r$  of the route r by U units. This will increase the capacity  $x_t$  of each single-link route  $t \in r$  by U units. Since  $i j \in r$ , the free capacity  $X_{ij}$  is increased by U units.
- 2. The physical link i j does not exist: Select one of the shortest routes  $r \in \mathcal{R}_{ij}$  such that each single-link route  $t \in r$  has unused capacity  $y_t > U$ . Increase the free capacity  $X_{ij}$  by U units by increasing the capacity  $x_r$  of route r by U units. This will decrease the capacity  $x_t$  of each single-link route  $t \in r$  by U units.



Figure 1 The core NSF network

#### 6. **RESULTS**

Consider the model [5] presented in figure 1 which is a fictitious representation of the core NSF ATM backbone consisting of eight nodes each connected by one bi-directional link apart from nodes 3 and 4 and 7 and 8 which are each connected by 2 links. The network carries 6 services: the bandwidth requirement of the 1st service is 1 unit and the bandwidth requirements of services 2 through 6 are 3, 4, 6, 24 and 40 units respectively. The transmission capacity of each link is 5624 units. The traffic intensity for service 3 is given in Table 1. For all other services, the values in Table 1 are multiplied by 0.5. The revenue earned by a carried call per unit of time is assumed to be independent of how the call is routed and is equal to the bandwidth of the call.

We first perform a robust test of the local reconfiguration algorithm: we design a VPCN from the physical network using only local reconfig-

nodes	1	2	3	4	5	6	7	8
1	_	6	7	1	9	5	2	3
2	7		<b>24</b>	3	31	15	6	9
3	8	25		4	<b>37</b>	18	7	11
4	1	3	3	-	4	7	1	1
5	11	33	39	5	_	<b>24</b>	9	15
6	5	14	16	2	21	-	4	6
7	2	<b>5</b>	6	1	8	4	-	2
8	3	8	10	1	12	6	2	

Table 1 Traffic intensity matrix

uration. This is done by initializing the simulator with the capacities of the physical links rather than with the capacities of the VPCs. We also provide the simulator with the routes which compose the VPCs as calculated by XFG: no capacities are initially assigned to the routes. In the second test we start with the fully-meshed VPCN constructed by XFG and use local reconfiguration to deal with random mismatches between offered traffics and VPC capacities. In all the simulations the threshold parameter  $D_{ij}$  is set to 12 and the reallocation unit U is set to 4. Each simulation run simulates 2 and a half million offered calls.

Table 3 presents the capacities of the VPCs resulting from these two tests. The column labeled VPCN presents the optimal VPCN constructed by XFG. Columns VPCN<sub>0</sub> and VPCN<sub>1</sub> present the VPCNs constructed by applying the local reconfiguration algorithm to the physical network and to the optimal VPCN respectively. Table 2 presents the network GoS for these two tests as well as for the optimal VPCN.

Tables 2 and 3 show that though the construction of a VPCN from a physical network is not the primary function of local reconfiguration, the near-optimal (in terms of the revenue rate) VPCN<sub>0</sub> constructed in this way suggests the utility of the local reconfiguration method. Perhaps the local algorithm can be used in situations where there are insufficient data to use XFG. Although the revenue rate generated by VPCN<sub>0</sub> is near optimal, the routes and the capacities assigned to the routes in VPCN<sub>0</sub> differ from those in the optimal VPCN. However, many near optimal VPCNs may exist.

The second test evaluates the local reconfiguration algorithm when it is used to deal with random mismatches between traffics and VPC capacities. Tables 2 and 3 show as expected (compare the columns labeled VPCN and VPCN<sub>1</sub>) that the local reconfiguration algorithm makes small changes to the optimal VPCN and the revenue rate derived from VPCN<sub>1</sub> remains close to the optimal value.

Table 4 compares local reconfiguration and dynamic alternate routing (DAR) when the traffic intensities to all O-D pairs in the network are increased by 20%. The DAR circuit reservation parameter is set

	VPCN	VPCN <sub>0</sub>	VPCN <sub>1</sub>
% blocking	1.7	2.6	1.6
% revenue lost	4.7	7.5	4.8
revenue rate	21277	20664	21265

Table 2 Global and local reconfiguration: network GoS

O-D	VPCN	$\mathrm{VPCN}_{0}$	$VPCN_1$	O-D	VPCN	$VPCN_0$	$VPCN_1$
1,2	2130	3220	2242	3,4	2451	2892	2435
1,3	808	844	860	3,5	3353	3228	3529
1,4	148	184	120	3,6	2544	2064	2544
1,5	793	1068	909	3,7	1282	880	1282
1,6	567	556	591	3,8	1809	2620	1661
1,7	263	280	258	4,5	439	296	467
1,8	366	288	342	$4,\!6$	859	752	827
2,3	2361	2176	2433	4,7	449	220	449
2,4	366	344	350	4,8	444	264	444
2,5	2599	2420	2303	$^{5,6}$	2023	1844	2023
2,6	1459	1380	1511	5,7	800	864	748
2,7	616	704	604	5,8	1241	1528	1269
2,8	903	1004	967	6,7	1296	3412	1252
				6,8	1226	800	1226
				7,8	3145	3904	3121

Table 3 Global and local reconfiguration: VPC capacities

to 4. The results labeled LOCAL show the effect of local reconfiguration; the results labeled NONE apply when local reconfiguration and alternative routing are disabled; the results labeled DAR apply when alternative routing is enabled and local reconfiguration is disabled. Alternative routing and local reconfiguration achieve almost the same overall blocking and revenue rate. Local reconfiguration achieves a better GoS for low bandwidth calls and far fewer low bandwidth calls are lost. In

Table 4 Local reconfiguration (LOCAL) versus dynamic alternate routing (DAR)

service class $s$ bandwidth $b_s$		all	1	2	3	4	5	6
			1	3	4	6	24	40
blocking	NONE	4.47	0.09	0.45	0.67	1.21	9.46	18.70
probability	LOCAL	3.77	0.01	0.02	0.06	0.24	8.54	17.42
	DAR	3.60	0.08	0.25	0.35	0.65	8.40	19.63
revenue	NONE	23567	326	976	2596	1937	7108	10624
rate	LOCAL	23848	327	980	2612	1956	7180	10793
	DAR	23324	316	946	2521	1885	7023	10633
attempts	LOCAL	9933	4834	1589	2358	827	203	122
	DAR	66163	9906	9716	19757	9786	8939	8059

particular table 4 shows that for the model under evaluation, local reconfiguration is attempted far fewer times that is alternative routing, with a corresponding reduction in signalling load. This is evident for the high bandwidth calls where alternative routing mostly fails. The revenue gain attributable to local reconfiguration is small. The rate of earning revenue for a class s call is proportional to the call's effective bandwidth requirement  $b_s$ : the improved GoS offered to low bandwidth calls therefore does not earn much revenue but will earn subscriber approval.

## 7. CONCLUSION

This paper is concerned with dynamic VPCN redesign as a resource management control in ATM networks where transmission capacity is reserved on the communication links in order to form dedicated logical paths for each origin-destination flow. We present a local reconfiguration algorithm designed to adapt an optimal VPCN to changes in the offered traffics. We present simulation results which indicate that, for the model under consideration, the local reconfiguration is robust, stable, and is capable of adapting the VPCN to reduce the call rejection count and thus limit the need for alternative routing.

#### References

- N. Anerousis and A.A. Lazar. Virtual Path Control for ATM Networks with Call Level Quality of Service Guarantees. *IEEE ACM Transactions on Networking* 6:2, April 1998, pp 222–236.
- [2] S.A. Berezner and A.E. Krzesinski. Call Admission and Routing in ATM Networks Based on Virtual Path Separation. IFIP TC6/WG6.2 Proceedings 4th International Conference on Broadband Communications, Stuttgart, Germany, April 1998. Chapman & Hall (Eds PJ Kühn and R Ulrich), pp 461–472.
- [3] R.J. Gibbens and F.P. Kelly. Dynamic routing in fully connected networks. IMA Journal of Mathematical Control and Information, Vol. 7 (1990) pp 77-111.
- [4] C. Lawrence, J.L. Zou and A.L. Tits. User's Guide for CFSQP Version 2.5. Report Number TR-94-16r1, Electrical Engineering Department and Institute for Systems Research, University of Maryland, College Park, MD 20742 USA.
- [5] D. Mitra, J.A. Morrison and K.G. Ramakrishnan. ATM Network Design and Optimization: a Multirate Loss Network Framework. Proceedings IEEE INFOCOM '96, pp 994–1003.